

Recovery of Protein Structure from Contact Maps

Michele Vendruscolo¹, Edo Kussell² and Eytan Domany¹

¹ *Department of Physics of Complex Systems, Weizmann Institute of Science, Rehovot 76100, Israel*

² *Department of Structural Biology, Weizmann Institute of Science, Rehovot 76100, Israel*

Background: Prediction of a protein's structure from its amino acid sequence is a key issue in molecular biology. While dynamics, performed in the space of two dimensional contact maps eases the necessary conformational search, it may also lead to maps that do not correspond to any real three dimensional structure. To remedy this, an efficient procedure is needed to reconstruct 3D conformations from their contact maps. Such a procedure is relevant also for interpretation of NMR and X-ray scattering experiments.

Results: We present an efficient algorithm to recover the three dimensional structure of a protein from its contact map representation. First we show that when a physically realizable map is used as target, our method generates a structure whose contact map is essentially similar to the target. Furthermore, the reconstructed and original structures are similar up to the resolution of the contact map representation. Next we use non-physical target maps, obtained by corrupting a physical one; in this case our method essentially recovers the underlying physical map and structure. Hence our algorithm will help to fold proteins, using dynamics in the space of contact maps. Finally we investigate the manner in which the quality of the recovered structure degrades when the number of contacts is reduced.

Conclusion: The procedure is capable of assigning quickly and reliably a 3D structure to a given contact map. It is well suited to be used in parallel with a dynamics in contact map space to project a contact map onto its closest physically allowed structural counterpart.

I. INTRODUCTION

Considerable effort has been devoted to finding ways to predict a protein's structure from its known amino acid sequence $\mathbf{A} = (a_1, a_2, \dots, a_N)$. The *contact map* of a protein is a particularly useful representation of its structure [1,2]. For a protein of N residues the contact map is an $N \times N$ matrix \mathbf{S} , whose elements are $S_{i,j} = 1$ if residues i and j are in contact and $S_{i,j} = 0$ otherwise. 'Contact' can be defined in various ways: for example, in a recent publication Mirny and Domany (MD) [3] defined contact $S_{i,j} = 1$ when a pair of heavy (all but hydrogen) atoms, one from amino acid i and one from j , whose distance is less than 4.5\AA , can be found. Secondary structures are easily detected from the contact map. Alpha helices appear as thick bands along the main diagonal, since they involve contacts between one amino acid and its four successors. The signature of parallel or anti-parallel beta sheets are thin bands, parallel or anti-parallel to the main diagonal. On the other hand, the overall tertiary structure is not easily discerned. The main idea of MD was to use this representation to perform a search, executed *in the space of possible contact maps* \mathbf{S} , for a fixed sequence \mathbf{A} , to identify maps of low "energy" $\mathcal{H}(\mathbf{A}, \mathbf{S})$. They defined the energy $\mathcal{H}(\mathbf{A}, \mathbf{S})$ as the negative logarithm of the probability that structures, whose contact map is \mathbf{S} , occur for a protein with the sequence \mathbf{A} ; therefore a map of low energy corresponds to a highly probable structure.

One of the most problematic aspects of their work was the fact that by performing an unconstrained search in the space of contact maps, i.e. freely flipping matrix elements from 1 to 0 and vice versa, one obtains maps of very low energy which have no physical meaning, since they do not correspond to realizable conformations of a polypeptide chain. To overcome this problem, MD introduced heuristic restrictions on the possible changes one is allowed to make to a contact map, arguing that if one starts with a physically realizable map, the moves allowed by these restrictive dynamic rules will generate maps that also are physically realizable. Even though their heuristic rules did seem to modify the dynamics in the desired way, there is no rigorous proof that indeed one is always left in the physical subspace, there is no clear evidence that the resulting

rules are not too restrictive and, finally, the need to start with a physical fold, copied from a protein of known structure, may bias the ensuing search and get it stuck in some local minimum of the energy.

The aim of the present publication is to present a method to overcome these difficulties. The idea is to provide a test, which can be performed "on line" and in parallel with the dynamics in the space of contact maps, which will "project" any map onto a nearby one which is guaranteed to be in the subspace of physically realizable maps. That is, for a given target contact map \mathbf{S} , we search for a conformation that a "string of beads" can take, such that the contact map \mathbf{S}' of our string is similar (or close) to \mathbf{S} . Needless to say, the contact map associated with a string of beads is, by definition, physical.

This particular aim highlights the difference between what we are trying to accomplish and existing methods [2,4–8], that use various forms of distance geometry [9] to predict a 3D structure of a real protein from its contact map. Rather than being concerned with obtaining a structure that is close to a real crystallographic one, we want mainly to check whether a map \mathbf{S} is physically possible or not, and if not - to propose some \mathbf{S}' which *is* physical and, at the same time, is not too different from \mathbf{S} . The method has to be fast enough to run in parallel with the search routine (that uses $\mathcal{H}(\mathbf{A}, \mathbf{S})$ to identify candidate maps \mathbf{S} of low energy). Another important requirement is to be able to recover contacts that do not belong to secondary structure elements and may be located far from the map's diagonal. Such contacts are important to nail down the elusive *global fold* of the protein. We believe that the main advantage of performing a dynamic search in the space of contact maps is the ease with which such contacts can be introduced, whereas creating them in a molecular dynamic or Monte Carlo procedure of a real polypeptide chain involves coherent moves of large sections of the molecule - moves which take a very long time to perform. To make sure that this advantage is preserved, our method must be able to efficiently find such conformations, if they are possible, once a new target contact has been proposed.

We are currently working on combining the method presented here with dynamics in the space of contact maps. The results of the combined procedure will be presented in a future publication.

The plan of the paper is as follows.

1. We give a detailed explanation of the method.
2. We show how it works on native maps of proteins with the number of residues ranging from $N=56$ to $N=581$. Success of the algorithm is measured in terms of the number of contacts recovered and the root mean square displacements of the recovered 3D structures from the native ones.
3. We study the answers given by our algorithm when it faces the task of finding a structure, using an unphysical contact map as its target. As the first test, we added and removed contacts at random to a physical map and found that the reconstructed structure did not change by much, i.e. we still could reconstruct the underlying physical structure. As a second test we used the constrained dynamic rules proposed by MD; starting from an experimental contact map, we obtained a new map by a denaturation-renaturation computer experiment. Since the MD rules are heuristic, this map is not guaranteed to be physical. Our reconstruction method projects a non-physical map onto one that is close to it *and* physically allowed.
4. We discuss the extent to which the quality of the structure, obtained from a map, gets degraded when the number of given contacts is reduced. This issue has considerable importance beyond the scope of our present study, since experimental data (disulfide bridge determination, crosslinking studies, NMR) often are available only for a small number of contacts. Clearly the more contacts one has the smaller is the number of possible conformations of a chain that are consistent with the constraints contained in the contact map. The question we address is when does this reduction of possible conformations suffice to define the corresponding structure with satisfactory accuracy.

II. METHOD

In this work we adopt a widely used definition of contact: two amino acids, a_i and a_j are in contact if their distance $d(a_i, a_j)$ is less than a certain threshold d_t . The distance $d(a_i, a_j)$ is defined as

$$d(a_i, a_j) = |\mathbf{r}_i - \mathbf{r}_j|, \quad (1)$$

where \mathbf{r}_i and \mathbf{r}_j are the coordinates of the C_α atoms of amino acids i and j .

The algorithm is divided into two parts. The first part, *growth*, consists of adding one monomer at a time, i.e. a step by step growth of the chain. The second part, *adaptation*, is a refinement of the structure, obtained as a result of the growth stage, by local moves. In both stages, to bias the dynamics, we introduce cost functions defined on the basis of the contact map. Such cost functions contain only geometric constraints, and do not resemble the true energetics of the polypeptide chain.

A. Growth

We first describe the growth stage of our procedure.

1. Single monomer addition. Here $N_t = 5$.

Suppose we have grown $i - 1$ monomers and we want to add point i to the chain. To place it, as shown in Fig. 1 we generate at random N_t trial positions, (typically $N_t = 10$),

$$\mathbf{r}_i^{(j)} = \mathbf{r}_{i-1} + \mathbf{r}^{(j)}, \quad (2)$$

where $j = 1, \dots, N_t$ and $\mathbf{r}^{(j)}$ is a vector whose direction is selected from a uniform distribution, whereas its length is distributed normally with average r_a and variance σ . Since in our representation monomers identify the C_α positions, we took $r_a = 3.79$ and $\sigma = 0.04$. We assign a probability $p^{(j)}$ to each trial in the following way. For each trial point $\mathbf{r}_i^{(j)}$ we calculate the contacts that it has (see eq. (1)) with the previously positioned points $\mathbf{r}_1, \dots, \mathbf{r}_{i-1}$. Contacts that should be present, according to the given contact map, are encouraged and contacts that should not be there are discouraged according to a cost function E_g that will be specified below. One out of the N_t trials is chosen according to the probability

$$p^{(j)} = \frac{e^{-E_g^{(j)}/T_g}}{Z}, \quad (3)$$

where the normalization factor is given by

$$Z = \sum_{j=1}^{N_t} e^{-E_g^{(j)}/T_g}. \quad (4)$$

The notation for the cost function E_g and for the parameter T_g that guide the growth are chosen in the spirit of the Rosenbluth method [10] to suggest their reminiscence to energy and temperature, respectively.

2. Chain growth.

The step by step growth presented in the previous section optimizes the position of successive amino acids along the sequence. The main difficulty in the present method is that the single step of the growing chain has no information on the contacts that should be realized many steps (or monomers) ahead. To solve this problem, we carry out several attempts (typically 10) to reconstruct the structure, choosing the best one. In practice, this is done as follows.

For each attempt, when position $\mathbf{r}^{(j)}$ is chosen for monomer i according to Eq. (3), its probability is accumulated in the weight

$$W_i = \prod_{k=1}^i p_k^{(j)} \quad (5)$$

When we have reached the end of the chain we store the weight W_N . The trial chain with the highest W_N is chosen.

3. Cost function.

The probabilities in Eq. (3) are calculated using the following cost function :

$$E_g^{(j)} = \sum_{k=1}^{i-1} f_g(r_{ik}^{(j)}) , \quad (6)$$

where $r_{ik}^{(j)} = |\mathbf{r}_i^{(j)} - \mathbf{r}_k|$, and

$$f_g(r_{ik}) = d \cdot a_g(S_{ik}) \cdot \vartheta(d_t - r_{ik}) , \quad (7)$$

The enhancing factor $d = i - k$ is introduced to guide the growth towards contacts that are long ranged along the chain; ϑ is the Heaviside step function and the constant a_g can take two values; $a_g(S_{ik} = 0) \geq 0$ and $a_g(S_{ik} = 1) \leq 0$. That is, when a contact is identified in the chain, i.e. $r_{ik} < d_t$, it is either "rewarded" (when the target map has a contact between i and k), or penalized. In this work we have grown chains with $a_g(0) = 0$. In this case, for a given contact map S , the function f_g only rewards those contacts that are realized and should be present. No cost is paid if contacts that are not in the map are realized by the chain (false positive contacts). Typically we chose the values $a_g(1) = -1.0$ and $T_g = 1$.

B. Adaptation

When we have grown the entire chain of N points, we refine the structure according to the following scheme. We choose a point i at random and try, using a crankshaft move, to displace it to \mathbf{r}'_i . We use a local cost function $E_a^{(i)}$:

$$E_a^{(i)} = \sum_{k=1}^{i-1} f_a(r'_{ik}) , \quad (8)$$

where $r'_{ik} = |\mathbf{r}'_i - \mathbf{r}_k|$, and

$$f_a(r_{ik}) = a_a(S_{ik}) \cdot \vartheta(d_t - r_{ik}) , \quad (9)$$

Note that the enhancing factor d has been omitted, so that f_a does not favor contacts between monomers that are distant along the chain. The displacement is accepted with probability π , according to the standard Metropolis prescription

$$\pi = \min(1, \exp(-\Delta E_a/T_a)) \quad (10)$$

where ΔE_a is the change in the cost function E_a induced by the move and T_a is a temperature-like parameter, used to control the acceptance ratio of the adaptation scheme.

A key ingredient of our method is annealing [11]. As in all annealing procedures, the temperature-like parameter T_a is decreased slowly during the simulation to help the system find the ground state in a rugged energy landscape.

In our method, however, instead of using simulation time as a control parameter on the temperature, we chose the number n of missing contacts. Two regimes were roughly distinguished. In the first regime many

contacts are missed and the map is very different from the target one. In the second regime few contacts are missed, and the map is close to the target. The parameters \mathbf{a}_a and T_a are interpolated smoothly between values suitable for these two limiting cases. In the first regime, we strongly favor the recovery of contacts that should be realized, whereas in the second regime we strongly disfavor contacts that are realized but should not be present. We set, as shown in Fig. 2

$$\mathbf{a}_a^{(n)}(S) = \mathbf{a}^f(S) + [\mathbf{a}^i(S) - \mathbf{a}^f(S)]\sigma(n) , \quad (11)$$

and

$$T_a^{(n)} = T_a^f + (T_a^i - T_a^f)\sigma(n) . \quad (12)$$

The function $\sigma(n)$ interpolates between the initial value \mathbf{a}^i and the final value \mathbf{a}^f ,

$$\sigma(n) = \frac{2}{1 + e^{-\alpha_g n}} - 1 . \quad (13)$$

By choosing \mathbf{a}^i , \mathbf{a}^f , T_a^i , T_a^f and α_g we define the two regimes, far from and close to the target map.

C. Chirality

A contact map contains no information about chirality. When an overall structure is reconstructed, the mirror image conformation is equally legitimate, having the same contact map. Since existing proteins do have a definite chirality, we are allowed to supply this information.

The $C_\alpha - C_\alpha$ distance constraints are loose enough to allow a local refinement of the reconstructed structure, with no loss in our geometrical cost function. Alpha helices can be detected as a thick band along the main diagonal of a contact map. A preliminary scan of the map identifies the sections that should be reconstructed as alpha helices. Next, we push the C_α 's in the alpha helices to positions that give the correct chirality, which is formally defined as the normalized triple product

$$c_i = \frac{\mathbf{v}_i \times \mathbf{v}_{i+1} \cdot \mathbf{v}_{i+2}}{|\mathbf{v}_i| \cdot |\mathbf{v}_{i+1}| \cdot |\mathbf{v}_{i+2}|} , \quad (14)$$

where $\mathbf{v}_i = \mathbf{r}_i - \mathbf{r}_{i-1}$.

In a typical alpha helix, $c_i = c_o = .778$ [6]. To refine the chirality of the preliminary chain that was obtained from the map by growth and adaptation as described above, we perform an additional Monte Carlo procedure. This procedure uses as "energy" a function that strongly favors the value quoted above for c :

$$E_c = a_c \left[\frac{2}{1 + \exp[-\alpha_c(c - c_o)^2]} - 1 \right] , \quad (15)$$

Since our Monte Carlo moves do not conserve the $C_\alpha - C_\alpha$ bond length, we added a term E_b to the energy function

$$E_b = a_b \left[\frac{2}{1 + \exp[-\alpha_b(r - r_a)^2]} - 1 \right] . \quad (16)$$

At each step a monomer i is selected randomly and its position displaced to

$$\mathbf{r}'_i = \mathbf{r}_i + \delta , \quad (17)$$

where δ is a small random vector. The total variation in the cost function, $E_a + E_c$, is evaluated with

$$c = \frac{c_{i-1} + c_i + c_{i+1}}{3} \quad (18)$$

used in Eq. (15).

Growth and adaptation yield a particular recovered structure, \mathbf{C} . We first create $\bar{\mathbf{C}}$, the mirror image of \mathbf{C} and use both structures as initial states for the final refinement procedure.

Usually either \mathbf{C} or $\bar{\mathbf{C}}$ evolves to a structure with the correct value of the average chirality rather quickly by our Monte Carlo process, while the mirror image does not, due to the lower compatibility of the latter structure with the correct chirality.

D. Remarks.

We devote the rest of this section to the discussion of alternative strategies that we have tried. Some of these may prove to be useful in future applications for more complex problems, but we have found that they are non necessary for the specific task dealt with in this work. We present these experiments since they underscore some non trivial aspects of the problem.

Adaptation alone: It is interesting to note that for short chains ($N < 200$) we can skip the growth stage; starting from a random structure, the adaptation procedure alone suffices to recover the correct set of contacts. The computer time needed for recovery increases, however, very fast with N . Since we are interested in recovering the structure in as short a time as possible, growth must be used, especially for long chains. We found that starting the adaptation stage from a grown (versus random) initial chain speeds up the procedure by a factor of about ten for proteins of length $N \simeq 100$. Moreover, for longer chains ($200 < N < 1000$), the cost function landscape is rougher and reconstruction by adaptation alone becomes unfeasible.

Piecewise growth: The importance of local contacts (i.e. contacts that involve amino acids nearby along the chain) versus non local ones has been discussed recently in the literature [12,13]. In these works evidence is given in support of the idea that non local interactions are decisive to stabilize the folded structure. A long standing alternative hypothesis [14] is that the folded structure is stabilized mainly by local interactions. We can test in the present work whether the *purely geometrical* (versus energetic) part of the reconstruction can or cannot be helped much by emphasizing the role of local contacts. To this end we used secondary structure elements as guidelines for the step by step growth. To implement this kind of growth instead of growing the entire chain of N amino acids, a section of M steps is built, with M ranging from 4 to 10 to match the size of a turn in an alpha helix or in a beta sheet. A set of sections is generated and the one with the best weight, according to Eq. (5), is chosen. Consistently with the findings in Refs. [12,13], we found that this secondary structure driven growth is not helping much the recovery.

A related idea is to optimize the relative positions of successive secondary structure elements. To realize this idea we have tried the following method. Similarly as above, sections of chain of M steps are grown, but now M is chosen randomly from 20 to 50. In this way we explore the space forward on the length scale of secondary structures to hook important contacts, i.e. those that fix the positions of secondary structures relative to each other. This scheme biases the growth to build a bridge to the next important contact, which usually is either inside a secondary structure or between different secondary structures. This method also allow to go back if too many mistakes are detected. As was the case for the previous attempted method, our experience suggests that this forward exploration is not necessary for solving the present task.

For multidomain proteins we tried to grow one domain at a time and then refining the structure by an adaptation cycle. The overall results were, however, only slightly affected.

An alternative idea that we tried is to bias the growth towards reaching a particular 'fixed point' [15]. For example, if it is known that two amino acids i and $j = i + k$ should be in contact then it is possible to bias the formation of a loop of length k . This method is well suited for very sparse contact maps, where it is easy to identify target points for the growing chain. We have verified that in dense maps the reconstruction speed is not increased by this scheme, due to the cumbersome identification of the target points.

Different cost functions: As mentioned above, we have used $a_g(0) = 0$. In general, this fact could lead to an overcompactification of the final structure. To assign unfavorable weight to false positive contacts, we should set $a_g(0) > 0$. This would introduce, however, frustration to the growth process, since it is guided by positive and negative energies. We discuss here the results of a possible way that we have tested to bias the growth away from conformations that contain spurious contacts that are not present in the map. We have assigned a positive cost $a_g(0) > 0$ to generating a spurious contact (i, j) if the closest existing contact (as measured on the map) is more than a distance of R units a way, e.g.

$$\min_{(h,k)} \sqrt{(i-h)^2 + (j-k)^2} > R, \quad (19)$$

where (h, k) run over all the existing contacts $S_{h,k} = 1$ in the given map. For the proteins we have analyzed (see Table I), we have extensively scanned possible values for R and $a_g(0)$. We found that that there is a strongly frustrated regime, for small R and large $a_g(0)$, where reconstruction is hindered; and a weakly

frustrated regime, for large R and small $a_g(0)$, where the efficiency of the reconstruction is only slightly improved with respect to using $a_g(0) = 0$. The intermediate regime, (typically $R = 5 - 10$ and $a_g(0) = 0.1 - 0.01$ for the values $a_g(1) = -1.0$ and $T_g = 1$ given above), may prove to be useful for proteins longer than those tested in the present work.

As for the functional form of the cost function, another possible choice, following Bohr *et al.* [5,6], is to smooth the step function that defines a contact with a sigmoid: Again, we did not find this necessary to achieve fast reconstruction.

In principle it would be possible to add to the function f_g of (7) a hard core repulsion

$$h(r) = \sigma_0(r - r_0)^{-\alpha} , \quad (20)$$

to try to overcome a general problem which arises when working with distance inequalities: an overcompaction of the globule, as measured, e.g. by the gyration radius. In practice, a good recovery prevents automatically the overlap between alpha carbons and the addition of such a term is not necessary.

III. EXPERIMENTAL CONTACT MAPS.

In this section we present results about the reconstruction of experimental contact maps as taken from PDB. Since our purpose, as explained in the Introduction, is to use the reconstruction in connection with dynamics, we chose $d_t = 9\text{\AA}$ to obtain the most faithful representation of the energy of the protein [3]. Such a threshold is determined by the requirement that the average number of $C_\alpha - C_\alpha$ contacts for each amino acid is roughly equal to the respective numbers obtained with the all-atom definition of contacts.

Two dissimilarity measures between structures are widely used. The most commonly adopted is the root mean square distance D

$$D = \sqrt{\frac{1}{N} \sum_{i=1}^N (\mathbf{r}_i - \mathbf{r}'_i)^2} , \quad (21)$$

where one structure is translated and rotated to get a minimal D . The standard procedure, described in Ref. [16], was used. Another possible choice is the distance D'

$$D' = \sqrt{\frac{1}{N^2} \sum_{i,j=1}^N (r_{i,j} - r'_{i,j})^2} , \quad (22)$$

The dissimilarity measure between contact maps is defined as the Hamming distance

$$D^{\text{map}} = \sum_{j>i} |S_{ij} - S'_{ij}| , \quad (23)$$

which counts the number of mismatches between maps \mathbf{S} and \mathbf{S}' .

For several proteins, we present in Fig. 3 the distances D and D' plotted vs the chain length N . The proteins considered, (with their respective lengths N and number of contacts N_c) are

The values of D and D' presented in Fig. 3 were obtained by averaging over 100 reconstruction runs for chains up to $N = 223$ and over 10 runs for longer chains. Error bars represent the variances as obtained from the corresponding sets of runs, as shown, for example, for the proteins 6pti bovine pancreatic trypsin inhibitor) and 1trmA (rat trypsin, chain A) in Fig. 4.

In Fig. 5 we show the contact map for the protein 6pti, $N = 56$, as taken from PDB, that was used as a target to construct a chain. The contact map of a typical reconstructed chain is also shown. In this particular case none of the 342 original contacts were missed and only two false positive contacts were added. These are close to clusters of correct contacts, indicating slight local differences with the crystallographic structure. The distances recorded in this case were $D' = 1.06$ and $D = 1.56$.

In Fig. 6 we show similar results for the larger protein 1trmA, with $N = 223$ and 1595 contacts. For clarity we have separated the experimental contact map from the reconstructed one. In the particular case shown there are 9 missing contacts and 84 false positives, and the corresponding distances are $D' = 1.34$ and $D = 1.59$. On average, in the 100 runs, 6 contacts were missed and 75 false positives were spuriously added. As in the case of 6pti, wrong contacts are mostly neighboring correct ones. Averages distances are $\langle D' \rangle = 1.3 \pm 0.1 \text{ \AA}$ and $\langle D \rangle = 1.6 \pm 0.2 \text{ \AA}$ (see also Fig. 4).

Using the distances D and D' to assess the quality of our results is misleading, since we are searching only for a chain that *reproduces the contacts* of a given map, whereas D and D' measure similarity between *structures*. Information, which is all-important to obtain low values for D and D' , such as the positions of amino acids that belong to loops, or slight rotations of secondary structures, is not contained at all in the map. For example, for the two-domain protein 1pii (phosphoribosylanthranilate isomerase), which has the largest distance in Fig. 3, only *two out of 3070 contacts were missed*, on the average, in the 10 reconstruction runs. However, changes in the relative orientations of the two domains lead to large distances. In fact, the target native maps were nearly perfectly reconstructed for all proteins tested.

We turn now to estimate the range of expected values for D and D' . The lower limit of our resolution for the chain, imposed by the geometrical constraints contained in the contact map, is about 1 \AA . To support this statement, we present the results of the following test. We subjected the native maps of 6pti, 1acx (actinoxanthin) and 1trmA from PDB and we subjected them to an adaptation cycle at low T_a . No native contacts were lost and no spurious ones were generated throughout the simulation, even though the structure (i.e. the positions of the beads) did vary; the most probable value for the distance D' between the generated structures was found to be around 1 \AA . This result clearly indicates the extent to which the contact map representation does not allow to nail down one specific structure to arbitrary precision. This ambiguity is compatible with the usual experimental resolution of PDB structures and hence the contact map representation is useful. Moreover, from low temperature flash photolysis experiments [17], X rays diffraction results analysis [18] and molecular dynamics simulations [19], the native fold of a protein is believed to consist of a set of conformational substates rather than of a unique structure [20].

The upper limit of the range of expected distances in our reconstruction is that between two completely unrelated structures, which can be as large as 15 \AA .

The conclusion of our studies is that our method produces, using a native contact map as target, a structure whose contact map is in nearly perfect agreement with the target. Furthermore, the distance of this reconstructed chain from the native structure is quite close to the resolution that can be obtained from the information contained in contact maps.

IV. NON PHYSICAL CONTACT MAPS

As stated in the Introduction, our main purpose is to develop a strategy to construct a three dimensional structure, starting from a given set of contacts, even if these contacts are not physical, i.e. not compatible with any conformation allowed by a chain's topology. In such a case we require our procedure to yield a chain whose conformation is as "close" as possible to the contact map we started with. The exact measure of such closeness depends on the source of the non-physicality, as will be demonstrated in two examples described below.

Our first examples of non physical contact maps were obtained by randomizing a native contact map; this was done by flipping M randomly chosen entries. Contacts between consecutive amino acids (neighbors along the chain) were conserved.

A typical contact map with noise is shown in Fig. 7. The protein is 6pti, whose map has 342 contacts. We show the native map and the target map obtained by flipping at random $M = 100$ entries of the native map, together with the map produced by our method. The most important conclusion that can be drawn from Fig. 7 is that isolated non-physical contacts are identified as such and ignored and the underlying physical contact map is recovered.

The dependence of this recovery on the noise level is shown in Fig. 8, where we present the average distance of the final structure from the uncorrupted 6pti map, for various values of M . Averages were taken over 10

reconstruction runs. The distance for totally unrelated structures for 6pti is around 8\AA . It is remarkable that up to $M < 200$ a fair resemblance to the experimental structure is still found.

The results of a similar study of a longer protein, 1trmA, with $M = 400$ flipped contacts, is shown in Fig. 9. For the particular case shown, distances to the PDB structure are $D' = 2.1\text{\AA}$ and $D = 2.4\text{\AA}$. The PDB map has 1595 contacts.

The family of possibly non physical contact maps, which is most relevant to our program, is produced by using the heuristic constrained dynamic rules in contact map space that were introduced by MD. Following them, we started with a native map of 6pti; when using a threshold $d_t = 8.5\text{\AA}$, the map derived from the PDB structure has 289 contacts, represented by open circles above the diagonal of Fig. 10. We have recomputed the energy parameters for the present definition of contacts (which involves C_α atoms only) and for a threshold of 8.5\AA [21]. The energy of the native 6pti map, obtained in this way, is 36.81. This contact map is subjected to repeated denaturation-renaturation cycles, using the constrained dynamics introduced by MD. We first heat the protein, inducing its unfolding, which is signaled by melting of secondary structure elements in the contact map. For moderate temperature shocks the protein is generally able to refold upon annealing [3].

In this work, we add a second step to this experiment, by subjecting the contact map obtained by constrained dynamics to our reconstruction procedure. To discuss in some detail the result of this combined scheme, we introduce three classes of contacts: we denote by A the contacts present in \mathbf{S}_A , the experimental contact map of protein 6pti (native contacts); by B the contacts present in \mathbf{S}_B , the contact map obtained by constrained dynamics in contact map space, starting from \mathbf{S}_A ; and by C the contacts in \mathbf{S}_C , the reconstructed contact map obtained using \mathbf{S}_B as a target. We present in Table II the total number of contacts in each class and the number of contact in common between two classes. Map \mathbf{S}_B has 255 contacts, 215 of which are in common with map \mathbf{S}_A . This difference is due to 74 missing contacts and 40 spurious ones (see also Fig. 10, above the diagonal). The energy of \mathbf{S}_B is 13.35. The reconstructed map \mathbf{S}_C has 310 contacts, 251 of which are in common with the target map \mathbf{S}_B . The difference arises from 4 missing contacts and 61 false positives. This reconstruction score is significantly larger than those typically obtained by applying directly our reconstruction procedure to native maps of proteins of similar size (see e.g. discussion about Fig. 5). This suggests, although without proving it, that the first step of the experiment, when we apply the rules of constrained dynamics introduced by MD is not guaranteed to yield a physically realizable map. However, from a closer inspection we can derive an overall consistency argument that implies that the MD rules do not drive the system very far away from the physical region in contact map space. The map \mathbf{S}_C and the native map \mathbf{S}_A , have 249 common contacts. \mathbf{S}_C has 40 missing and 61 false positives with respect to \mathbf{S}_A . Nevertheless, the distances in the 3D structures are $D' = 2.14$ and $D = 2.97$ respectively, indicating a rather successful refolding.

From these results we argue that MD rules alone are possibly not enough restrictive to keep the trajectory of the system in the physical region of the contact map space during a denaturation-renaturation experiment. This problem can be corrected by the reconstruction procedure discussed in this work. Our procedure projects the contact map obtained by MD rules onto a contact map which is admissible by construction. Rather consistently, the projected map is quite close to the target one.

V. REDUCING THE NUMBER OF CONTACTS

In this section we address a very important issue; the effect of reducing the number of contacts on the accuracy of the reconstructed structure. Even though resolving this problem is not essential for our goals, its resolution is an interesting spinoff obtained from our algorithm. The question is relevant to a number of problem areas where contacts are of importance, such as protein structure determination from NMR data [22], studies of DNA and crosslinked polymers. The latter are known to undergo a vulcanization transition from a liquid phase to a frozen amorphous phase if the number of contacts exceeds a critical value [23]. The stochastic reconstruction method described in this work is rather general and potentially applicable to these systems as well.

In real proteins the number of contact scales with the length N of the chains, as shown in Fig. 11 for a representative set of 246 proteins taken from PDB. Fitting the data with a single power law

$$M = aN^\nu, \quad (24)$$

yields best fit for $\nu = 1.07$ but the data are also compatible with linear scaling (e.g. $\nu = 1$) as well as with a combination of linear scaling and surface corrections

$$M = aN + bN^{2/3}. \quad (25)$$

All three fits are shown in the figure and are nearly indistinguishable on the scale used. Fig. 11 was obtained using the definition of contacts as given by Eq. (1) with a threshold of $d_t = 9\text{\AA}$. In the range $5\text{\AA} < d_t < 9\text{\AA}$, only the prefactor a changes, while the exponent ν remains the same. This result holds also for the MD definition of a contact. It has been proposed that in order to have a compact structure, the minimum number of contacts of a random heteropolymer should scale linearly with N [24,25] or with $N/\ln N$ [26]. These findings suggest that in proteins the number of contacts required to determine the native fold also cannot scale with a power that is much less than linear with N . The relevant issue, to which considerable effort has been recently devoted [7,8], concerns how small can be the prefactor a , in order to achieve reasonable reconstruction of protein structure from incomplete experimental contact maps. We address this point by analyzing the feasibility of the reconstruction as the threshold d_t is decreased. The smaller d_t , the smaller is the number of contacts present in the contact map.

We now present detailed studies of the protein lacx with $N = 108$. For $d_t = 9\text{\AA}$ the number of contacts was 652; for $d_t = 6$ this number becomes 253 and 154 for $d_t = 5$. Note that the optimal parameters used for annealing depend on d_t ; for $d_t = 5\text{\AA}$, for example the values $a^f(1) = -5$, $a^i(1) = -20$, $a^f(0) = 0.1$ and $a^i(0) = 0.5$ were used.

In all cases our method produced chains whose contact maps were in nearly perfect agreement with the respective target maps (deviating by one or two spurious contacts). The distances of the corresponding structures from the native one are, however, very different as d_t decreases. As shown on Fig. 12, the values of the average distances D', D (obtained from 100 runs for each d_t) increase, from $1 - 2\text{\AA}$ for $d_t = 9\text{\AA}$ to $5 - 8\text{\AA}$ for $d_t = 5\text{\AA}$. This striking increase of D with decreasing d_t shows that even when the target contact map is essentially perfectly recovered, the corresponding structure can be very different from the true one. This suggests that for low values of d_t more information than what is contained in the contact map should be provided to get acceptable resemblance to the experimental structure.

VI. SUMMARY AND CONCLUSIONS

We presented a stochastic method to derive a 3D structure from a contact map representation. We have shown that for physically realizable target maps our method is very fast and reliable to find a chain conformation whose contact map is nearly identical to the target. Moreover, the method is able to find a good candidate structure even when the target map has been corrupted with non-physical contacts.

The information contained in the contact map derived from an experimental structure allows us to reconstruct a conformation, which is relatively close to that of the original structure, as already observed by Havel *et al.* [2]. There is, however, an intrinsic limit in the resolution of a contact map. When a threshold of 9\AA between C_α atoms is used to define contact, the distance between two typical structures, that are both compatible with the contact map, is of the order of 1\AA . Some recent work focused on the identification of supplemental information needed to improve the resolution of the reconstruction [5–8]. Usually, energy terms are introduced to favor secondary structures, side chain burial, hydrogen bonding and sulfur bridges. An example for such terms is the chirality term used in the present work. Although these issues are of considerable importance for experimental structure determination, we consider the problem from a different perspective.

The contact map representation is intimately connected with the pairwise parametrization of the energy. We are currently studying the dynamics in contact map space, controlled by such a pairwise energy function. We believe that such a study will reveal whether the contact map representation, together with the assumptions implicit in working with a pairwise, contact-based approximation for the energy, suffice to single out the native state of a protein.

This research was supported by grants from the Minerva Foundation, the Germany-Israel Science Foundation (GIF) and by a grant from the Israeli Ministry of Science. E. D. thanks the members of the Laboratoire de Physique Statistique at the Ecole Normale Supérieure, Paris for their hospitality. E.K. thanks the K. Kupcinet fellowship program for partial support of his stay as a summer student at the Weizmann Institute.

-
- [1] Lifson, S., & Sander, C. (1979). Antiparallel and parallel beta strands differ in amino acids residue preferences. *Nature* **282**, 109-111.
 - [2] Havel, T. F., Crippen, G. M., & Kuntz, I. D. (1979). Effect of distance constraints on macromolecular conformation. II. Simulation of experimental results and theoretical predictions. *Biopolymers* **18**, 73-81.
 - [3] Mirny, L., & Domany, E. (1996). Protein fold recognition and dynamics in the space of contact maps. *Proteins: Struct. Funct. and Gen.* **26**, 391-410.
 - [4] Havel, T. F., & Wuthrich, K. (1985). An evaluation of the combined use of nuclear magnetic resonance and distance geometry for the determination of protein conformation in solution. *J. Mol. Biol.* **182**, 281-294.
 - [5] Bohr, J., *et al.*, & Petersen, E. F. (1993). Protein structures from distance inequalities. *J. Mol. Biol.* **231**, 861-869.
 - [6] Lund, O., Hansen, J., Brunak, S., & Bohr, J. (1996). Relationship between protein structure and geometrical constraints. *Protein Science* **5**, 2217-2225.
 - [7] Aszodi, A., & Taylor, W. R. (1996). Homology modelling by distance geometry. *Folding & Design* **1**, 325-334.
 - [8] Skolnick, J., Kolinski, A., & Ortiz, A. R. (1997). MONSSTER: A method for folding globular proteins with a small number of distance restraints. *J. Mol. Biol.* **265**, 217-241.
 - [9] Crippen, G., & Havel, T. F. (1988). *Distance geometry and molecular conformation*. Wiley, New York.
 - [10] Rosenbluth, M. N., & Rosenbluth, A. W. (1955). *J. Chem. Phys.* **23**, 356.
 - [11] Kirkpatrick, S., Gelatt, C. D. Jr, & Vecchi, M. P. (1983). Optimization by simulated annealing. *Science* **220**, 671-680.
 - [12] Abkevich, V. I., Gutin, A. M., & Shakhnovich, E. I. (1995). Impact of local and non-local interactions on thermodynamics and kinetics of protein folding. *J. Mol. Biol.* **252**, 460-471.
 - [13] Govindarajan, S. & Goldstein, R. (1995). Why are some protein structures so common?. *Proc. Natl. Acad. Sci. USA* **93**, 3341-3345.
 - [14] Anfinsen, C. & Scheraga, H. (1975). Experimental and theoretical aspects of protein folding. *Advan. Protein Chem.* **29**, 205-300.
 - [15] Vendruscolo, M. (1997). Modified configurational bias Monte Carlo for simulation of polymer systems. *J. Chem. Phys.* **106**, 2970-2976.
 - [16] Kabsch, W. (1976). A solution for the best rotation to relate two sets of vectors. *Acta Cryst.* **A32**, 922-923; Kabsch, W. (1978). A discussion of the solution for the best rotation to relate two sets of vectors. *Acta Cryst.* **A34**, 827-828.
 - [17] Austin, R. H., Beeson, K. W., Eisenstein, L., Frauenfelder, H., & Gunsalus, I. C. (1975). Dynamics of ligand binding to myoglobin. *Biochemistry* **14**, 5355.
 - [18] Frauenfelder, H., Parak, F., & Young, R. D. (1988). Conformational substates in proteins. *Ann. Rev. Biophys. Biophys. Chem.* **17**, 451-479.
 - [19] Elber, R. & Karplus, M. (1987). Multiple conformational states of proteins: Molecular dynamics of myoglobin. *Science* **235**, 318.
 - [20] Frauenfelder, H., Sligar, S., & Wolynes, P. G. (1991). The energy landscapes and motions in proteins. *Science* **254**, 1598-1603.
 - [21] Energy parameters are available upon request.
 - [22] Wuthrich, K. (1986). *NMR of Proteins and Nuclear Acids*. Wiley, New York.
 - [23] Goldbart, P. M., & Zippelius, A. (1993). Amorphous solid state of vulcanized macromolecules: a variational approach. *Phys. Rev. Lett.* **71**, 2256; Castillo, H. E., Goldbart, P. M., & Zippelius, A. (1994). Distribution of localization lengths in randomly crosslinked macromolecular networks. *Europhys. Lett.* **28**, 519.
 - [24] Camacho, C. J. (1996). Entropic barriers, frustration and order: basic ingredients in protein folding. *Phys. Rev. Lett.* **77**, 2324-2328; Camacho, C. J. & Schanke, T. (1997). From collapse to freezing in random heteropolymers. *Europhys. Lett.* **37**, 603-608.

- [25] Gutin, A. M. & Shakhnovich, E. I. (1994). Statistical mechanics of polymers with distance constraints. J. Chem. Phys. **100**, 5290-5293.
- [26] Bryngelson, J. D. & Thirumalai, D. Internal constraints induce localization in an isolated polymer molecule. Phys. Rev. Lett. **76**, 542-546.

FIG. 1. Single monomer addition.

FIG. 2. Annealing functions for the parameters used in adaptation. See Eqs. (11) and (12).

FIG. 3. Average distances $\langle D \rangle$ (lower) and $\langle D' \rangle$ (upper) vs chain length N for the proteins listed in Table I.

FIG. 4. Histograms of the distances D' (dots) and D (open circles) for the 100 runs used to test the reconstruction procedure. Data are presented for proteins 6pti (left) and 1trmA (right).

FIG. 5. Contact maps for protein 6pti for a threshold $d_t = 9\text{\AA}$. Dots are the PDB data, open circles the output of the reconstruction procedure. None of the target contacts is missed and two spurious ones are added.

FIG. 6. Contact map of protein 1trmA. Experimental contact map (above diagonal) and reconstructed one (below diagonal).

FIG. 7. Above diagonal: target map (crosses) obtained by randomizing the underlying physical map (dots) of the protein 6pti. Below diagonal: reconstructed contact map (open circles) obtained by using the noise-corrupted map as target. The physical map is also shown (dots).

FIG. 8. Average distances D (open circles) and D' (dots) vs noise M for 6pti.

FIG. 9. Above diagonal: reference map (crosses) obtained by randomizing the underlying physical map (dots) of protein 1trmA. Below diagonal: reconstructed contact map (open circles) obtained using the noise-corrupted map as target.


FIG. 10. Reconstruction of a contact map obtained by the constrained dynamics introduced by MD. Above diagonal: reference map (dots), obtained by a denaturation-renaturation cycle starting from the experimental map of protein 6pti (open circles). Below diagonal: reconstructed contact map (crosses), obtained by using as a target the contact map (dots) that was obtained by constrained dynamics.


FIG. 11. Scaling of the number M of contacts with the length N of the proteins. Data refer to 246 proteins taken from PDB, and to a threshold $d_t = 9\text{\AA}$.



FIG. 12. Average distances $\langle D' \rangle$ (dots) and $\langle D \rangle$ (open circles) as a function of the threshold d_t for protein 1acx.


TABLE I. List of PDB proteins used to test the reconstruction procedure.


TABLE II. Number of contacts in the classes A, B and C explained in the text. Also the number of common contacts between classes is reported.
































protein	N	N_c	protein	N	N_c
6pti	56	342	2sodO	151	1066
2ci2	65	350	1hge1	185	1084
1tlk	103	610	1akeA	214	1348
5cytR	103	644	1trmA	223	1595
1ltsD	103	571	1abe	305	2179
9rnt	104	623	1pii	452	3070
1acx	108	652	3gly	470	3383
2trxA	108	628	3cox	500	3680
1f3g	150	1049	1gal	581	4369
1aak	150	922			

A	B	C	AB	AC	BC
289	255	310	215	249	251